independent rules, yet medical decision making is a complex process in which many factors are interrelated. Thus, attempting to represent medical decision-making as a discrete set of independent rules, no matter how complex, is a task that can, at best, result in a first order approximation of the process. This places an inherent limitation on the quality of feedback that can be provided. As a consequence it is extremely difficult to develop feedback that explicitly takes into account all information available on the patient. One might speculate that the lack of widespread acceptance of such systems may be due to the fact that their recommendations are often rejected by physicians. These systems must be made more valid if they are to enjoy widespread acceptance among physicians.

The proposed MENTOR system is designed to address the significant problem of adverse drug reactions by means of a computer-based monitoring and feedback system to influence physician decision-making. It will employ principles of artificial intelligence to create a more valid system for evaluating therapeutic decision-making.

The work in the MENTOR project is intended to be a collaboration between Dr. Blaschke at Stanford and Dr. Speedie at the University of Maryland. Dr. Speedie provides the expertise in the area of artificial intelligence programming. Dr. Blaschke provides the medical expertise. The blend of previous experience, medical knowledge, computer science knowledge and evaluation design expertise they represent is vital to the successful completion of the activities in the MENTOR project.

*C. Highlights of Research Progress*

The MENTOR project was initiated in December 1983. The project has been funded by the National Center for Health Services Research since January 1, 1985. Initial effort has focused on exploration of the problem of designing the MENTOR system. Work has begun on constructing a system for monitoring potassium in patients with drug therapy that can adversely affect potassium. Antibiotics, dosing in the presence of renal failure, and digoxin dosing have been identified as additional topics of interest.

*E. Funding Support*

Title: MENTOR: Monitoring Drug Therapy for Hospitalized Patients

Principal Investigators:
Terrence F. Blaschke, M.D.
Division of Clinical Pharmacology
Department of Medicine
Stanford University

Stuart M. Speedie, Ph.D.
School of Pharmacy
University of Maryland

Funding Agency: National Center for Health Services Research

Grant Identification Number: 1 R18 HS05263

Total Award: January 1, 1985 - December 31, 1988 $485,134 Total Direct Costs

Current Period: January 1, 1986 - December 31, 1986 $182,820 Total Direct Costs

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

### *A. Medical Collaborations and Program Dissemination via SUMEX*

This project represents a collaboration between faculty at Stanford University Medical Center and the University of Maryland School of Pharmacy in exploring computer-based monitoring of drug therapy. SUMEX, through its communications capabilities, facilitates this collaboration of geographically separated project participants by allowing development work on a central machine resource and file exchange between sites.

### *B. Sharing and Interactions with Other SUMEX-AIM Projects*

Interactions with other SUMEX-AIM projects has been on an informal basis. Personal contacts have been made with individuals working on the ONCOCIN project concerning system development issues. Dr. Perry Miller has also been of assistance by providing software for advisory generation. Given the geographic separation of the investigators, the ability to exchange mail and programs via the SUMEX system as well as communicate with other SUMEX-AIM projects is vital to the success of the project.

### *C. Critique of Resource Management*

To date, the resources of SUMEX have been fully adequate for the needs of this project. The staff have been most helpful with any problems we have had and we are quite satisfied with the current resource management.

## III. RESEARCH PLANS

### *A. Project Goals and Plans*

The MENTOR project has the following goals:

1. Implement a prototype computer system to continuously monitor patient drug therapy in a hospital setting. This will be an expert system that will use a modular, frame-oriented form of medical knowledge, a separate inference engine for applying the knowledge to specific situations, and automated collection of data from hospital information systems to produce therapeutic advisories.

2. Select a small number of important and frequently occurring medical settings (e.g., combination therapy with cardiac glycosides and diuretics) that can lead to therapeutic misadventures, construct a comprehensive medical knowledge base necessary to detect these situations using the information typically found in a computerized hospital information system and generate timely advisories intended to alter behavior and avoid preventable drug reactions.

3. Design and begin to implement an evaluation of the impact of the prototype MENTOR system on physicians' therapeutic decision-making as well as on outcome measures related to patient health and costs of care.

1986 will be spent on prototype development in four content areas, design and implementation of the basic knowledge representation and reasoning mechanisms, and preliminary interfacing to existing patient information systems.

### *B. Justification and Requirements for Continued SUMEX Use*

This project needs continued use of the SUMEX facilities for two reasons. First, it provides access to an environment specifically designed for the development of AI systems. The MENTOR project focuses on the development of such a system for drug monitoring that will explore some neglected aspects of AI in medicine. This environment is necessary for the timely development of a well-designed and efficient MENTOR system. Second, access to SUMEX is necessary to support the collaborative efforts of geographically separated development teams at Stanford and the University of Maryland. Furthermore, the MENTOR project is predicated on the access to the SUMEX resource free of charge over the next two years. Given the current restrictions on funding, the scope of the project would have to be greatly reduced if there were charges for use of SUMEX.

*C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM*

A major long-range goal of the MENTOR project is to implement this system on a independent hardware system of suitable architecture. It is recognized that the full monitoring system will require a large patient data base as well as a sizeable medical knowledge base and must operate on a close to real-time basis. Ultimately, the SUMEX facilities will not be suitable for these applications. Thus we intend to transport the prototype system to a dedicated hardware system that can fully support the the planned system and which can be integrated into the SUMC Hospital Information System. For this purpose a VAX 750 and two XEROX 1186 workstations have been acquired and our development efforts are gradually being transferred to them.

*D. Recommendations for Future Community and Resource Development*

In the brief time we have been associated with SUMEX, we have been generally pleased with the facilities and services. However, it is clearly evident that the users' almost insatiable demands for CPU cycles and disk space cannot be met by a single central machine. The best strategy would appear to be one of emphasizing powerful workstations or relatively small, multi-user machines linked together in a nationwide network with SUMEX serving as the its central hub. This would give the individual users much more control over the resources available for their needs, yet at the same time allow for the communications among users that have been one of SUMEX's strong points.

For such a network to be successful, further work needs to be done in improving the network capabilities of SUMEX to encourage users at sites other that Stanford. Further work is also needed in the area of personal workstations to link them to such a network. Given the successful completion of this work, it would be reasonable to consider the gradual phase-out of the central SUMEX machine over two or three years and its replacement by an efficient, high-speed communications server.

# IV.B.4. Rutgers Research Resource

**Rutgers Research Resource--Artificial Intelligence in Medicine**

**Principal Investigators:**
**Casimir Kulikowski, Sholom Weiss**
**Rutgers University, New Brunswick, New Jersey**

## I. SUMMARY OF RESEARCH PROGRAM

### *A. Goals and Approach*

The fundamental objective of the Rutgers Resource is to develop a computer based framework for advancing research in the biomedical sciences and for the application of research results to the solution of important problems in health care. The central concept is to introduce advanced methods of computer science - particularly in artificial intelligence - into specific areas of biomedical inquiry. The computer is used as an integral part of the inquiry process, both for the development and organization of knowledge in a domain and for its utilization in problem solving and in processes of experimentation and theory formation.

An essential part of the resource is directed to methodological problems of knowledge representation and to the development of computer-based systems for acquiring, managing, and improving knowledge bases, and for constructing expert reasoning models in medicine. Equally fundamental are the problems of how best to use knowledge bases and models in processes of interpretation/diagnosis, planning, theory formation, simulation, and effective man-machine communication. These are problems we are studying in the Resource in the context of several system building efforts that address themselves to specific tasks of clinical decision-making and model development and testing.

Resource activities include research projects (collaborative research and core research) training/dissemination projects, and computing services in support of user projects.

### *B. Medical Relevance and Collaborations*

In 1985-86 we continued the development of several versatile systems for building and testing consultation models in biomedicine. The EXPERT system has had many of its capabilities enhanced in the course of collaborative research in the areas of rheumatology, ophthalmology, and clinical pathology.

In *ophthalmology* we have developed a knowledge representation scheme for treatment planning which is both natural and efficient for encoding the strategies for choosing among competing and cooperating treatment plans. This involves a ranking of treatments according to their characteristics and desired effects as well as contraindications. A diagnosis and treatment planning program for ocular herpes was developed using this scheme. Our main collaboration continues to be with Dr. Chandler Dawson of the Proctor Foundation, UCSF.

In *rheumatology*, the model for rheumatological diseases now includes detailed diagnostic criteria for 26 major diseases. The management advice and treatment planning has been developed further. The Resource researchers have developed new representational elements for EXPERT in response to the needs of the rheumatology research Politakis originally developed a coordinated system called SEEK (System for

Empirical Experimentation with Expert Knowledge) which provides interactive assistance to the human expert in testing, refining and updating a knowledge base against a data base of trial cases. A generalized version of SEEK, SEEK2, has been developed during the past year. Dr. Lindberg of the National Library of Medicine, and Dr. Sharp of the University of Missouri are the project leaders in developing the rheumatology knowledge base for this effort.

In *clinical pathology* our main collaboration has been with Dr. Robert Galen (Cleveland Clinic Foundation), with whom we have developed the serum protein electrophoresis model which is incorporated into an instrument with a scanning densitometer. This instrument with interpretive reporting capabilities has now been on the market for over a year, and is located at several hundred clinical sites. We are making good progress developing a knowledge based system for the interpretation of CPK/LDH isoenzymes.

In biomedical modeling applications we are experimenting with several prototype models for giving advice on the interpretation of experimental results in the field of enzyme kinetics, in conjunction with Dr. David Garfinkel. His PENNZYME program has been linked to a model in EXPERT, which allows the user to interpret the progress of the model analysis, and a framework for the design of experiments in this domain has been formulated.

*C. Highlights of Research Progress*

Research has continued on problems of representation, inference and control in expert systems. Emphasis has been placed this year on problems of knowledge base acquisition, empirical testing and refinement of reasoning (the SEEK2 system). From a technological point of view, the market availability of the interpretive reporting version of a scanning densitometer, and the development of models for eye care consultation that run on microprocessor systems (Apple IIe, IBM-PC) represents an important achievement for AIM research in showing its practical impact in medical applications.

- Knowledge Base Refinement: SEEK is a system which has been developed to give interactive advice about rule refinement during the design of an expert system. The advice takes the form of suggestions for possible experiments in generalizing and specializing rules in an expert model that has been specified based on reasoning rules cited by a human expert. Case experience, in the form of stored cases with known conclusions, is used to interactively guide the expert in refining the rules of a model. The design framework of SEEK consists of a tabular model for expressing expert-modeled rules and a general consultation system for applying a model to specific cases. This approach has proven particularly valuable in assisting the expert in domains where the logic for discriminating two diagnoses is difficult to specify, and we have benefited primarily from experience in building the consultation system in rheumatology. During the past year, a newer SEEK2 system has been developed that has enhanced capabilities, including a more generalized knowledge base and an automatic pilot capability to proceed with knowledge base refinements. The original work on SEEK was recognized by the international HUSPI award for medical expert system research.

- Technology Transfer: Important technology transfer milestones have also been achieved this year: the instrument interpretation EXPERT program for serum protein has been widely disseminated, as has the Ocular Herpes Treatment Program.

*D. Up-to-Date List of Publications*

The following is an update of publications in the Rutgers Resource for the period 1984 and 1985 (only publications not listed in previous SUMEX annual reports are presented here).

1. Apte, C. and Weiss, S., An Approach to Expert Control of Interactive Software Systems, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 7 No. 5, pp. 586-591, (1985).

2. Ginsberg, A., Weiss, S., and Politakis, P., SEEK2: A Generalized Approach to Automatic Knowledge Base Refinement, *Proceedings of the 1985 International Joint Conference on Artificial Intelligence* pp. 367-374 (1985).

3. Ostroff, J., Dawson, C., Kastner, J., Weiss, S., Kulikowski, C., Kern, K., An Expert Advisory System for Primary Eye Care in Developing Countries, *Proceedings 1985 Expert Systems in Government Conference*, pp. 490-495 (1985).

4. Ostroff, J., Dawson, C., Kastner, J., Weiss, S., Kulikowski, C., Kern, K., Preliminary Results from Field Testing of an Expert Advisory System for Primary Eye Care in Developing Countries, Proceedings AI-86 Conference, Long Beach, CA. (1986).

5. Galen, R. and Weiss, S., The CPK/LDH Diagnostic Program, *Proceedings of the 5th International Meeting on Clinical and Laboratory Organization and Management*, Haifa, in press (1985).

6. Weiss, S.M. and Kulikowski, C.A.: *A Practical Guide to Designing Expert Systems*, Rowman and Allanheld, 1984.

7. Kastner, J., Weiss, S., Kulikowski, C., and Dawson, C.: *Therapy Selection in an Expert Medical Consultation System for Ocular Herpes Simplex* Computers in Biology and Medicine, Vol. 14, No. 3, pp. 285-301 (1984).

8. Dawson, C., Kastner, J., Weiss, S., Kulikowski, C.: *A Computer-based Method to Provide Subspecialist Expertise on the Management of Herpes Simplex Infections of the Eye*, Proceedings International Symposium On Herpetic Eye Diseases, Belgium (1984).

9. Galen, R. and Weiss, S.: *Predictive Value Calculator*, American Society of Clinical Pathologists, Clinical Chemistry # CC 84-4 (1984).

10. Kastner, J., Dawson, C., Weiss, S., Kern, K., Kulikowski, C.: *An Expert Consultation System for Frontline Health Workers in Primary Eye Care*, Journal of Medical Systems, Vol. 8, No. 5 (1984).

Indicate by an asterisk (*) that the resource was given credit.

*E. Funding Support*

The Rutgers Research Resource on Artificial Intelligence in Medicine is funded under grant RR 02230-01 from the Division of Research Resources, Biotechnology Resources Program. Principal Investigators are Casimir A. Kulikowski, Professor and Chairman of the Department of Computer Science, and Sholom M. Weiss, Associate Research Professor of Computer Science.

The total direct costs for the period 1983-87 is $3,198,075.

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

### *A. Medical Collaborations and Dissemination*

The SUMEX-AIM facility provides a backup node where some of our medical collaborators can access programs developed at Rutgers. The bulk of the medical collaborative work outlined in I.B. above is centered at the Rutgers facility (the Rutgers-AIM node).

Dissemination activities continue to be an important responsibility of the Rutgers Resource within the AIM community. The following activities took place in the last year:

1. AIM Workshop (1985):

   Organized by Dr. Kingsland of the National Library of Medicine. It consisted of a series of presentations on AIM research and related work by members of the AIM community.

2. 1984 Hawaii International Conference On Systems Sciences:

   Dr. Weiss presented a paper on the expert system for front-line health workers, and Dr. Kulikowski chaired a session on knowledge based medical systems.

### *B. National AIM Projects at Rutgers*

The national AIM projects, approved by the AIM Executive Committee, that are associated with the Rutgers-AIM node are the following:

1. Attending Project, directed by Dr. Perry Miller of the Yale Medical Center, is doing much of the research on critiquing a physician's plan of management at Rutgers.

2. Medical Knowledge Representation project, headed by Dr. Chandrasekaran from Ohio State University, is doing most of its research on the Rutgers system.

3. Biomedical Modeling, by Dr. Garfinkel from the University of Pennsylvania.

4. INTERNIST/CADUCEUS project, headed by Dr. Myers and Dr. Pople from the University of Pittsburgh, has been using the Rutgers Resource as a backup system for development and experimentation.

### *C. Critique of SUMEX-AIM Resource Management*

Rutgers is currently using the SUMEX DEC-20 system primarily for communication with other researchers in the AIM community and with SUMEX staff, and also for backup computing in demonstrations, conferences and site visits. Our usage is currently running at less than 50 connect hours per year at SUMEX, with an overall connect/CPU ratio of about 30.

Rutgers is beginning to place more emphasis on the use of personal computers, and on network support needed to make these effective. SUMEX has been of significant help in their developmental efforts in networking workstation software.

## III. RESEARCH PLANS

### *A. Project Goals and Plans*

We are planning to continue along the main lines of research that we have established in the Resource to date. Our medical collaborations will continue with emphasis on development of expert consultation systems in rheumatology, ophthalmology and clinical pathology. The basic AI issues of representation, inference and planning will continue to receive attention. Our core work will continue with emphasis on further development of generalized expert system frameworks and also on AI studies in representations and problems of knowledge and expertise acquisition. We propose to work on a number of technology transfer experiments to micro processing that will be affordable by our biomedical research and clinical collaborators. We also plan to continue our participation in AIM dissemination and training activities as well as our contribution -- via the Rutgers computers -- to the shared computing facilities of the national AIM network.

### *B. Justification and Requirements for Continued SUMEX Use*

Continued access to SUMEX is needed for:

1. Backup for demos, and so forth.
2. Programs developed to serve the National AIM Community should be runnable on both facilities.
3. There should be joint development activities between the staffs at Rutgers and SUMEX in order to ensure portability, share the load, and provide a wider variety of inputs for developments.

### *C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM*

Our computing needs are based on a centralized computing resource accessible to distant users, and local workstations. We will continue to use SUMEX for backup purposes.

### *D. Recommendations for Future Community and Resource Development*

Use of personal computers and workstations is continuing to grow in the AIM community. We find that the biggest challenge is supporting these systems. Although some central computing will continue to be needed for communication and coordination, we believe that over the next few years all AIM research projects and even individual collaborators will come to have their own hardware. However, many of these community members (particularly the collaborators) will not be in a position to support hardware or software on their own. We would certainly expect SUMEX to continue to provide expert advice in this area. However, we believe it would be helpful for SUMEX to have a formal program to support smaller computers in the field. We envision this as including at least the following items:

- A central source of information on hardware and software that is likely to be of interest to the AIM community. SUMEX might want to become a distribution point for certain of this software, and even help coordinate quantity purchase of hardware if this proves useful.
- Assistance in support of hardware and software in the field. Depending upon the hardware involved, this might involve advice over the telephone or actual board-swapping by mail.

# IV.B.5. SOLVER Project

**SOLVER: Problem Solving Expertise**

**Dr. P. E. Johnson**
**Center for Research in Human Learning**
**University of Minnesota**

**Dr. James R. Slagle**
**Department of Computer Science**
**University of Minnesota**

**Dr. W. B. Thompson**
**Department of Computer Science**
**University of Minnesota**

## I. SUMMARY OF RESEARCH PROGRAM

### *A. Project Rationale*

This project focuses upon the development of strategies for discovering and documenting the knowledge and skill of expert problem solvers. In the last several years, considerable progress has been made in synthesizing the expertise required for solving extremely complex problems. Computer programs exist with competency comparable to human experts in diverse areas ranging from the analysis of mass spectrograms and nuclear magnetic resonance (Dendral) to the diagnosis of certain infectious diseases (Mycin).

Design of an expert system for a particular task domain usually involves the interaction of two distinct groups of individuals, "knowledge engineers," who are primarily concerned with the specification and implementation of formal problem solving techniques, and "experts" (in the relevant problem area) who provide factual and heuristic information of use for the problem solving task under consideration. Typically, the knowledge engineer consults with one or more experts and decides on a particular representational structure and inference strategy. Next, "units" of factual information are specified; that is, properties of the problem domain are decomposed into a set of manageable elements suitable for processing by the inference operations. Once this organization has been established, major efforts are required to refine representations and acquire factual knowledge organized in an appropriate form. Substantial research problems exist in developing more effective representations, improving the inference process, and in finding better means of acquiring information from either experts or the problem area itself.

Programs currently exist for empirical investigation of some of these questions for a particular problem domain (e.g. AGE, UNITS, RLL). These tools allow the investigation of alternate organizations, inference strategies, and rule bases in an efficient manner. What is still lacking, however, is a theoretical framework capable of reducing dependence on the expert's intuition or on near exhaustive testing of possible organizations. Despite their successes, there seems to be a consensus that expert systems could be better than they are. Most expert systems embody only the limited amount of expertise that individuals are able to report in a particular, constrained language (e.g. production rules). If current systems are approximately as good as human experts, given that they represent only a portion of what individual human experts know, then improvement in the "knowledge capturing" process should lead to systems with considerably better performance.

In order to obtain a broad view of the nature of human expertise, the SOLVER project includes studies in a variety of complex problem solving domains in addition to medicine. These include law, auditing, business management, plant pathology, and expert system design. We have observed that despite the apparent dissimilarities in these problem solving areas, there is reason to believe that there are underlying principles of expertise which apply broadly. Our project seeks to investigate these principles and to create tools to make use of that knowledge in practical expert systems.

*B. Medical Relevance and Collaboration*

Much of our research has been and will continue to be directly focused on medical AI problems. Galen, our experimental expert system in pediatric cardiology, is achieving expert levels of performance. Dr. Connelly is leading a project that is developing an expert system based platelet transfusion therapy monitoring program. Dr. Spackman is completing a doctoral thesis on the automated acquisition of rule knowledge in medical microbiology.

Some of our research has focused on problems in diagnostic reasoning and expertise in domains other than medicine. However, our experience indicates that principles of expertise and relevant knowledge engineering tools can cut across task domains. Galen is demonstrably a useful expert system implementation tool designed in the medical diagnostic task domain. Developments from our work in other domains affecting problems such as automated knowledge acquisition through rule induction and reasoning by analogy will have medical relevance.

Collaboration with Dr. James Moller in the Department of Pediatrics, Dr. Donald Connelly in the Department of Laboratory Medicine, at the University of Minnesota. Dr. Connelly has become a SUMEX user and is teaching a course in medical informatics. He has also initiated a project to create an expert system in platelet transfusion therapy.

Collaboration with Dr. Eugene Rich and Dr. Terry Crowson at St. Paul Ramsey Medical Center.

Dr. Kent Spackman is a post-doctoral fellow in medical informatics who is completing a Ph.D. thesis in Artificial Intelligence. Dr. Spackman is a resident at the University of Minnesota Hospitals and collaborates with the SOLVER project.

Recent research by Prof. Slagle has addressed the issue of economical information acquisition in medical diagnostic expert systems. A demonstration system has been implemented studying cases from the Program on the Surgical Control of the Hyperlipidemias. This expert system evaluates the results of treadmill ECG examinations.

*C. Highlights of Research Progress*

Accomplishments of This Past Year --

Dr. Slagle has published the results of his work on the design of the Merit system. Merit enables an expert system to direct the acquisition of information by finding questions with a high ratio of probable importance to difficulty. Merit has been incorporated in a general purpose expert system shell called GENIE (Generalized Inference Engine). In addition, Dr. Slagle has described an expert system that can be used to support the data collection and evaluation problems in a large clinical study: the Program on the Surgical Control of Hyperlipidemias (POSCH). POSCH is a long term randomized clinical trial with 838 post-myocardial infarction patients.

Dr. Connelly has supervised the implementation of two medical expert systems. One deals with *detection of deviations in time series by the human observer* and is an expert system to assist observers in the monitoring of laboratory data. The other is a *knowledge based system for improving transfusion practice.* The object of this system is to promote nearly optimal usage of platelet concentrates in a tertiary care setting through the use of an expert system.

Dr. Spackman has developed an algorithm useful for inductive inference in generating classification rules from empirical examples. The algorithm has been tested with data from several medical domains, including microbiology, endocrinology and pediatric cardiology.

Dr. Johnson has been engaged in continuing studies of medical expertise. Current research has included the study of memory processes of diagnosticians and the role of model-based reasoning in diagnosis. In addition, a project studying the role of context in analogical reasoning is under way, and a prototype inference engine using context directed reasoning by analogy has been implemented.

Drs. Thompson and Johnson have also undertaken a study of explanation facilities in expert systems. Explanation is related to the line of reasoning of the system and the needs of the users.

Research in Progress --

*Application of Galen in different problem domains*

One project to test the extensibility of Galen into other domains has recently been completed in an area of accounting (internal audit controls). The objective of the study was to formulate and test a model of the processes employed by audit managers and partners in reviewing and evaluating internal accounting controls. In collaboration with an international accounting firm, a rule base was developed and successfully tested against the actions of expert auditors.

Another project explored the extension of the Galen architecture into a problem in agriculture. The main objective of this research was to determine whether the basic postulates about expert reasoning made in Galen would apply successfully in another diagnostic domain. The problem domain chosen for this purpose was plant pathology. The system is currently going through evaluation and fine tuning to bring it up to an expert performance level. This system may be useful in the Extension Service at the University of Minnesota, which provides diagnostic information to farmers over the phone lines.

*Inference engine mechanisms*

In another area of research, a small rule based program called PROTEUS was recently completed for the problem of computer hardware diagnosis. Several hardware diagnostic programs are currently based on structures that are essentially equivalent to decision trees. As such a program's competence becomes larger, however, the program becomes increasingly difficult to modify. The purpose of the PROTEUS project was to investigate how a rule based system could be used to represent the same diagnostic knowledge as a decision tree and yet be easier to understand and modify due to the modular nature of production rules. PROTEUS represents the state of the problem being solved as a tree of symbols. For example, some portions of the tree represent lines of questioning that should be followed to get information about the problem. Other portions represent facts that have been obtained as a result of previous questions. Still others represent conclusions that were drawn from the these facts. PROTEUS works by using rewrite rules to transform the tree. A rewrite rule recognizes a subtree that matches the pattern in its antecedent and either replaces or modifies this subtree with a new tree from its consequent. This single general mechanism is used to ask questions, draw conclusions from the answers, and modify the line of questioning accordingly.

*Merit system for question selection*

Another direction in our current research has focused on developing a generalized inference engine as an extension and improvement of the ideas found in BATTLE, an expert consultant system for weapon to target allocation. The BATTLE inference engine models the relevant considerations for allocating a set of weapons to a set of targets. The system incorporates variants of several important techniques of artificial intelligence and makes the first use of the Merit system for question selection. The Merit system enables the system to direct the acquisition of information by finding questions with a high ratio of probable importance to difficulty. In the first phase of the allocation process, the system uses a computation network to determine the effectiveness of each individual weapon against each prospective target. The network, built by a domain expert in advance, allows reasoning with logical, Bayesian, and expert-defined operators. After the calculation of individual effectiveness values, a portion of the allocation tree is built to determine good allocation plans for the set of weapons. The individual effectiveness values are used to direct the traversal and pruning of the allocation tree.

This approach has been extended in the GENIE (GENeralized Inference Engine) system. GENIE was designed to provide greater representational power, simplified knowledge specification, and smaller storage requirements. GENIE can manipulate elementary concepts involving several objects at once and can represent the subset-superset relationships of the types of objects. The elementary concepts may have values taken from any well-defined data type. Moreover, representational structures such as nodes and edges are shared whenever possible, and some redundancies of the previous system are reduced, resulting in an increase of storage efficiency in GENIE. The current version of GENIE includes the use of the Merit questioning system and a menu-driven interface. GENIE not only represents a significant achievement in research on inference engines, but also provides as excellent tool for future research on Merit.

*Detection of deviations in time series by the human observer*

Surveillance and early detection of deviation from a homeostatic state are goals common to health care programs for the apparently healthy as well as for groups of patients known to have or have had specific diseases. Automated approaches to detecting deviations have the advantage of being reliably applied, traceable, consistent in outcome, and conserving of professional resources. However, these techniques may be relatively insensitive, require a moderately large number of observations before they become functional, and are not necessarily tuned to the detection of the specific type of deviation of clinical significance. The human observer using graphical trend displays coupled with specific domain knowledge may be a more sensitive and reliable detector of deviations at an earlier time in the course of clinical observations. But the human observer may be inconsistent, unavailable, distracted or expensive. By embodying the human rules of deviation detection in an expert system, perhaps the advantages of the human expert can be gained while at the same time the disadvantages are eliminated. A study has been done to test these hypotheses. Time trend graphs representing monthly monitoring of serum carcinoembryonic antigen (CEA) levels in simulated patients with surgically-removed breast cancer were presented to six human observers and a time series analysis based on a homeostatic model. Three of the human observers described their rationale in assigning a level of suspicion regarding the presence of an important deviation as the observation points on the graphs were serially revealed to them. The verbalizations were analyzed to develop rules for an expert system. Based on a small number of cases, the false positive rate of all three approaches (i.e., analytic method, observers, and expert system) appeared comparable. Observers detected deviations earlier than the analytic method in 13 of 39 instances and tied in 15. The expert system tied the observers in 19 of 38 instances and surpassed them seven times. This preliminary study suggests that human observers using time trend graphs can supplement analytic techniques in the early detection of deviations in biologic time series and that the human skill of graph reading can at least partially be extracted and embodied in an expert system.

*Knowledge based system for improving transfusion practice*

The goal of this project was to promote more nearly optimal usage of platelet concentrates in a tertiary care setting. Through the use of an expert system operating as an information gathering aid to blood bank personnel, professional knowledge and judgment will consistently, automatically and economically be brought to bear on each request for platelet concentrates. The approach used should lead to more consistent and complete professional overview and coordination of blood product use, the development of general and specific guidelines for appropriate use of platelet concentrates, decreased blood bank administrative costs related to blood product use review, and decreased costs to patients. Finally, the project will demonstrate an approach to medical care quality assurance that more realistically deals with the complexity, variability and breadth of clinical problems than is possible with conventional criteria list techniques.

The specific aims of the research are:

1. to develop a set of transfusion practice guidelines for platelet concentrates and to employ these guidelines in an active tertiary care setting, to develop a supplementary guideline set in the form of a knowledge base allowing disease-specific, therapy-specific, trend and other contingencies to be taken into account when considering the appropriateness of transfusion, to develop an inexpensive computer-based monitoring system using expert system techniques to semi-automatically evaluate adherence to these transfusion guidelines,

2. to develop an automatically maintained platelet use database for providing physicians with practice pattern feedback, and to demonstrate the effect of this approach on platelet concentrate use and administrative costs.

At the present time, the platelet transfusion knowledge base acquired so far has been incorporated into an expert system implemented in TLC LISP running on an HP Vectra microcomputer. Automatic acquisition of laboratory data has been accomplished. Work is continuing in knowledge acquisition and computer to computer communication capability.

*Inductive Inference for Medical Diagnosis*

An inductive inference algorithm has been developed for generating classification rules from empiric examples. Such rules may be useful for the conceptual analysis of data and for the construction and refinement of knowledge bases for expert systems. This research addresses some deficiencies of currently available methods for deriving classification rules from empiric data, and explores the use of unate Boolean functions as a language for logical classification rules.

The goals of this research are to find computational methods for deriving comprehensible classification rules from empiric data, and to evaluate the performance of these methods in producing formal criteria for classification. Further, this research attempts to find and develop methods for creating rules (concept descriptions) which use a format known as "criteria tables" or "counting rules."

Such criteria tables may be viewed more abstractly as Boolean functions. In addition, they fall into a category known as unate Boolean functions, which may be defined as the set of logical expressions that can be formed using the connectives AND and OR, without the NOT connective. The algorithm that has been developed takes advantage of some of the interesting and useful properties of unate functions.

Currently available methods in machine learning, pattern recognition, and statistical discriminant analysis have several deficiencies that limit their usefulness in generating descriptions of concepts or in building expert systems. The pattern recognition and discriminant analysis methods fail to produce human-oriented, comprehensible results, and they do not take into account background knowledge about the problem domain. Machine learning algorithms (such as Michalski's AQ or Quinlan's ID3) produce logic-based results, but they have other drawbacks. For example, some cannot deal with inconsistencies, incomplete data or noise in the training examples; most produce discriminants which incorrectly overfit the data. Because they use languages with essentially unlimited expressiveness, they place no bound on the complexity of the derived rules, leading to complicated rules and a high degree of instability of the rules with varying training sets. Most provide for little or no incorporation of background knowledge.

To address these deficiencies, a logic-based approach to inductive inference is combined with background knowledge in an algorithm that allows appropriate handling of inconsistencies and noise, producing concept descriptions that can be expressed as unate Boolean functions or may be further simplified to criteria tables. A Boolean function is unate if no variable appears both complemented and uncomplemented in its minimum sum-of-products (disjunctive normal form) representation. From an intuitive point of view, unateness means that an individual piece of evidence (a variable or proposition) either contributes for or against the conclusion of the rule, but never reverses its influence when combined with other evidence.

The algorithm that has been developed accepts both binary and non-binary variables as input, and then creates propositional variables from the non-binary ones using techniques related to Michalski's AQ algorithm. It also accepts user-supplied domain and task knowledge which it uses to derive rules that more closely reflect the background knowledge and goals of humans. For example, the diagnostic performance goal of the rule can be specified. Possible goals may be 1) to minimize the expected mis-classification error (or cost), 2) to maximize the sensitivity, or 3) to maximize the specificity of the rule. The algorithm then handles conflicting examples (examples with identical features but belonging to different classes) by finding the unate function that best satisfies the given goal of the rule to be derived.

Criteria tables derived by the algorithm may contain an arbitrarily large number of categories of findings (e.g. major, minor, required, exclusions, etc.) but when the number of categories becomes too large, the comprehensibility of the result diminishes. Methods for creating a hierarchy or network of intermediate rules may help to improve comprehensibility in such cases, and developing such methods is a long-term goal of this research.

The algorithm has been tested with data from several medical domains, including microbiology, endocrinology and pediatric cardiology. The results from microbiology show that unate functions are adequate for expressing rules for differentiating groups of Enterobacteriaceae according to a panel of biochemical tests. The results from endocrinology indicate that machine learning adds new information above and beyond what is derivable using linear discriminant analysis. The results from pediatric cardiology reveal that most of the logical competitor sets of the Galen system may be differentiated equally well in thousands of different ways (i.e., with thousands of different criteria tables), adding confirmation to the impression that each expert diagnostician has his own individual strategies that appear to be equally valid.

An unanswered question is whether unate concepts are more natural or more easy to learn than conjunctive concepts or disjunctive concepts. The problem becomes more difficult to answer when one considers the tendency to use general rules of thumb modified by exceptions, because the empirical exceptions may violate the unateness condition and thus make it uncertain whether a concept is inherently unate or whether the exceptions necessarily make the concept disjunctive. Nonetheless, the comprehensibility of unate functions seems to be related to their correspondence to psychological models of the way people categorize.

*D. List of Relevant Publications*


1. Bailey, A. et al.: "Auditing, artificial intelligence, and expert systems," *Decision Support Systems: Theory and application*, A.B. Whinston (ed.), North-Holland Publications, 1985.

2. Connelly, D. and Johnson, P.E.: *Medical problem solving*. Human Pathology, 11(5):412-419, 1980.

3. Elstein, A., Gorry, A., Johnson, P. and Kassirer, J.: *Proposed Research Efforts*. In D.C. Connelly, E. Benson and D. Burke (Eds.), CLINICAL DECISION MAKING AND LABORATORY USE. University of Minnesota Press, 1982, pp. 327-334.

4. Feltovich, P.J.: *Knowledge based components of expertise in medical diagnosis*. Learning Research and Development Center Technical Report PDS-2, University of Pittsburgh, September, 1981.

5. Feltovich, P.J., Johnson, P.E., Moller, J.H. and Swanson, D.B.: *The Role and Development of Medical Knowledge in Diagnostic Expertise*. In W. Clancey and E.H. Shortliffe (Eds.), READINGS IN MEDICAL AI, Addison-Wesley, 1984, pp. 275-319.

6. Johnson, P.E.: *Problem Solving*. In ENCYCLOPEDIA OF SCIENCE AND TECHNOLOGY, McGraw-Hill, 1985.

7. Johnson, P.E., Moen, J.B., and Thompson, W.B.: *Garden Path Errors in Medical Diagnosis*. IN Bloc, L. and Coombs, M.J. (Eds.), COMPUTER EXPERT SYSTEMS, Springer-Verlag (in press).

8. Johnson, P.E.: *Cognitive Models of Medical Problem Solvers*. IN D.C. Connelly, E. Benson, D. Burke (Eds.), CLINICAL DECISION MAKING AND LABORATORY USE. University of Minnesota Press, 1982, pp. 39-51.

9. Johnson, P.E.: *What kind of expert should a system be?* Journal of Medicine and Philosophy, 8:77-97, 1983.

10. Johnson, P.E.: *The Expert Mind: A New Challenge for the Information Scientist* IN Th. M. A. Bemelmans (Ed.), INFORMATION SYSTEM DEVELOPMENT FOR ORGANIZATIONAL EFFECTIVENESS, Elsevier Science Publishers B. V. (North-Holland), 1984.

11. Johnson, P.E., Severance, D.G. and Feltovich, P.J.: *Design of decision support systems in medicine: Rationale and principles from the analysis of physician expertise*. Proc. Twelfth Hawaii International Conference on System Science, Western Periodicals Co. 3:105-118, 1979.

12. Johnson, P.E., Duran, A., Hassebrock, F., Moller, J., Prietula, M., Feltovich, P. and Swanson, D.: *Expertise and error in diagnostic reasoning*. Cognitive Science 5:235-283, 1981.

13. Johnson, P.E. and Hassebrock, F.: *Validating Computer Simulation Models of Expert Reasoning.* IN R. Trappl (Ed.), CYBERNETICS AND SYSTEMS RESEARCH. North-Holland Publishing Co., 1982.

14. Johnson, P.E. and Thompson, W.B.: *Strolling down the garden path: Detection and recovery from error in expert problem solving.* Proc. Seventh IJCAI, Vancouver, B.C., August, 1981, pp. 214-217.

15. Johnson, P.E., Hassebrock, F. and Moller, J.H.: *Multimethod study of clinical judgement.* Organizational Behavior and Human Performance 30:201-230, 1982.

16. Moller, J.H., Bass, G.M., Jr. and Johnson, P.E.: *New techniques in the construction of patient management problems.* Medical Education 15:150-153, 1981.

17. Sedlmeyer, R.L., Thompson, W.B. and Johnson, P.E.: *Knowledge-based fault localization in debugging.* The Journal of Systems and Software, vol. 3, no. 4 (Dec 83) pp. 301-307, Elsevier.

18. Sedlmeyer, R.L., Thompson, W.B. and Johnson, P.E.: *Diagnostic reasoning in software fault localization.* Proc. Eighth IJCAI, Karlsruhe, West Germany, August, 1983.

19. Slagle, J. et al.: "Expert systems in medical studies -- A new twist," *Proceedings of the Conference on Applications of Artificial Intelligence*, SPIE, 1986.

20. Slagle, J.: "An expert system for a resource allocation problem," *Communications of the ACM*, September, 1985.

21. Slagle, J.: "Alpha Beta Pruning," *Encyclopedia of artificial intelligence*, S. Shapiro (ed.), John Wiley and Sons, Inc., New York, 1985.

22. Slagle, J. and Gaynor, M.: "An intelligent control strategy for computer consultation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, March, 1984.

23. Smith, K.A., Farm, B., Johnson, P.E.: *Surface: A prototype expert system for selecting surface analysis techniques.* Proceedings of IEEE Conference on Computers and Comm., 1985.

24. Swanson, D.B.: *Computer simulation of expert problem solving in medical diagnosis.* Unpublished Ph.D. dissertation, University of Minnesota, 1978.

25. Swanson, D.B., Feltovich, P.J. and Johnson, P.E.: *Psychological Analysis of Physician Expertise: Implications for The Design of Decision Support Systems.* In D.B. Shires and H. Wold (Eds.), MEDINFO77, North-Holland Publishing Co., Amsterdam, 1977, pp. 161-164.

26. Thompson, W.B., Johnson, P.E. and Moen, J.B.: *Recognition-based diagnostic reasoning.* Proc. Eighth IJCAI, Karlsruhe, West Germany, August, 1983.

*E. Funding and Support*

Work on the SOLVER project is currently supported by a grant from the Control Data Corporation to Paul Johnson ($95,000; 1986-88) and by a grant from the Microelectronics and Information Sciences Center at the University of Minnesota to

Paul Johnson, William Thompson, James Slagle (Dept. of Computer Science), Harry Wechsler (Electrical Engineering), and Albert Yonas (Institute for Child Development) ($300,000; 1985-86).

Research in medical informatics is supported, in part, by a training grant from the National Library of Medicine, LM-00160, in the amount of $712,573 for the period 1984-1989. Dr. Connelly and Prof. Johnson are participants in this grant. The post-doctoral fellowship of Dr. Spackman is funded by this grant.

Dwan Family Fund, University of Minnesota Medical School, $6,000 (1985) to Paul Johnson for research assistant funding on the Galen project.

## II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

### *A. Medical Collaborations and Program Dissemination via SUMEX*

Work in medical diagnosis is carried out with the cooperation of faculty and students in the University of Minnesota Medical School and St. Paul Ramsey Medical Center.

The Galen system is now available from the University of Minnesota as an unsupported research tool for the study of recognition based reasoning systems.

### *B. Sharing and Interactions with Other SUMEX-AIM Projects*

Dr. Paul Johnson participated in a panel discussion with Dr. Harry Pople at the International Conference on Information Systems, Indianapolis, Ind., Dec. 1985. The panel addressed issues in the design and implementation of practical expert systems and expert system research.

### *C. Critique of Resource Management* (None)

## III. RESEARCH PLANS

### *A. Project Goals and Plans*

*Model based diagnostic reasoning*

Since human experts are notoriously poor at describing their own knowledge, one of our planned research objectives is to investigate problem solving tasks through which experts can reveal criteria for initiating specific hypotheses and methods for investigating those hypotheses. We have recently begun preliminary work into how symbolic models of a physical system may be used to acquire knowledge for diagnosing the cause of its observed faults. Previous work at Minnesota on both the Galen and PROTEUS projects (see above) suggests that a model based approach could be useful in acquiring diagnostic knowledge from experts as well as providing a framework for building an expert system that puts this knowledge to use.

We plan to investigate a model seen as a network of component frames that are connected by bidirectional communication links. Each frame represents a real or imaginary component of the system that is being studied. Each link represents a possible dependency between the states of the two components that it connects. For example, a model of a computer might include a memory frame, an arithmetic unit frame, and a bus frame. The arithmetic unit frame might be linked to the bus frame, and the bus frame linked to the memory frame. Symbolically represented information can flow back and forth between frames along links. Other frames that are not part of the computer itself may be included as well. For instance, it may be useful to imagine that all frames are linked to a general environment frame that describes conditions under which the computer is operating.

A component frame has an internal state that is represented symbolically within it. It also has a set of rules that specify relationships between the internal state and the information flowing in and out through the communications links. These rules operate bidirectionally. If a symptom is presented to a frame through one of its links, the frame's rules deduce an internal state that could plausibly cause the symptom. Alternatively, if the internal state changes so that a symptom would be expected, the rules work in the other direction and deduce a symptom that is sent out along a link. The model therefore functions much like a constraint network.

Once the physical model is built, it can be used as the heart of a diagnostic expert system. To use the model in diagnosis, descriptions of symptoms are placed in the internal state of its environment frame. Rules in this frame are applied to suggest changes in the internal states of other frames by sending results along their common communication links. These changes propagate through the model in the same way. Finally the rules in some frame deduce possible faults that could explain the observed symptoms. Each possible fault also causes conjectured symptoms to be deduced. Checks for the presence or absence of these symptoms could be used to confirm or disconfirm the fault. Results of these checks could then be fed back into the model so that incorrect conjectured faults can be removed and correct ones can be strengthened.

One question addressed in this work is what we call the *"inversion problem"* . While it is relatively easy to obtain knowledge from experts about what faults cause what symptoms, it is more difficult to obtain knowledge about what symptoms allow the diagnosis of faults. It is tempting to suppose that the first kind of knowledge, determining symptoms given faults, could be obtained and *inverted* to obtain the more diagnostically useful knowledge about determining faults given symptoms. This is difficult because one symptom may be caused by several different faults, so new knowledge must be provided to decide between them. Such information includes, but is not limited to: What level of detail and abstraction is appropriate for diagnostic reasoning about a system? What faults are the most plausible causes of an individual symptom? What procedures or tests can be executed to investigate a conjectured fault? How can the results suggest what to do next?

*Context directed analogical reasoning*

A second area of planned work is the investigation of reasoning by analogy within the context of law. Lawyers reason using legal rules, and they reason about the formation and foundation of the rules themselves. While current expert systems technology supports rule-based reasoning when the underlying rules are fixed, the processes involved in reasoning about rules is less well understood. An important feature of legal reasoning is the use of precedents. Precedents are used to discover legal rules and to formulate legal arguments about such rules. We plan to develop a computer model of reasoning with precedents that will be used to investigate the role of "context" in selection and use of precedents. The context selection mechanism is hypothesized to be represented by a system like the Galen system. The process of reasoning with precedents is an elaboration of the model proposed by Winston. The goal of this project is to develop a system that can simulate the reasoning processes of an attorney in selecting relevant precedents from a "knowledge-base" of case law.

*Machine learning*

A third area of planned work is machine learning. One research project currently underway is developing methods for the automated generation of comprehensible decision rules from empiric data, with emphasis on logic-based knowledge representation formats and on problems drawn from the domain of medicine. This work builds on some of the machine learning methodologies developed at the University of Illinois by R. S. Michalski and others.

This work addresses two shortcomings of previous work on induction of classification rules. These are: first, lack of comprehensibility of the induced rules; and second, lack of flexibility in specifying the diagnostic performance (sensitivity, specificity, or efficiency) desired for the rules that are to be derived.

Comprehensibility of the derived rules or descriptions can be enhanced by imposing restrictions on the format that the rules may take. For example, the restriction of rules to a unate Boolean function format follows the induction of rules that can often be simplified to a *"criteria table"* type of representation. The type of diagnostic performance a rule must have will depend on its purpose, and specifying the purpose may allow inductive inference algorithms to trade off small decrements in diagnostic performance for large increments in comprehensibility, or to increase their robustness in the face of noisy or uncertain data.

Successful development of these techniques will lead to enhanced capabilities for deriving rule bases for expert classification systems from empiric data, and will provide new methods for the conceptual analysis of data.

Preliminary results have been obtained for the problem of deriving rules for the identification of bacteria based on their biochemical profiles in the medical microbiology lab. Other problem domains under investigation are the analysis and interpretation of endocrine laboratory tests, and the induction of rules for the diagnosis of congenital heart disease, for comparison with the rules used in Galen.

Related research is also under way to determine methods of augmenting the process of acquiring the knowledge used to build rule bases for expert classification systems. The premises of this research are that efficient knowledge acquisition depends on the choice of representation, and that in classification domains the concept of *inconsistency* plays an important role. The objective of this research is to design, implement, and test a representation of the beliefs of an *"ideally rational agent"* that is capable of explicitly representing, automatically detecting, and using logical inconsistencies in the agent's set of beliefs. These inconsistencies form the basis for rules that generate class hypotheses, discriminate among them, reject invalid hypotheses, detect noisy or bad input data, and that can direct the search for unknown facts by a questioning facility.

We plan to continue development of a representational formalism for classification knowledge bases that enables inconsistencies to be manipulated as described above. This system is based on an existing nonmonotonic logic belief called *autoepistemic logic* (as defined by Robert Moore) and modified to include the principles of *relevance logic* (as defined by Anderson and Belnap). Having thus defined a representation for the knowledge base, we plan to develop a method for instantiating its concepts within a given application domain. Preliminary research results suggest that an efficient process for learning a rule base consists of first adding the expectations for each class into the logic base and then introducing test cases to the system one at a time. Because the system can automatically detect inconsistencies and determine their type and significance, it can suggest both that a rule is required and what its type should be. We expect the resulting rule base to be more focused and efficient, and to more closely duplicate the lines of reasoning of domain experts. The rule acquisitions tools will be tested in two different applications domains. The resulting expert systems will be evaluated for correctness of classification and similarity of their lines of reasoning with those of human experts.

*Intelligent query generation*

A fourth area of planned work will be on improving and extending the Merit control strategy described earlier. This is a new computer strategy for deciding what piece of evidence to look for at each stage of the decision process. Every proposition of interest is assigned a weight (called a merit value) proportional to its ability to alter the decision. At each stage of the consultation, the user is asked the question having the highest merit value. An efficient algorithm has been devised to find the proposition have the highest merit. After the relevant evidence has been gathered, the algorithm updates merit values to reflect the new information. The advantage of this approach is that in a time-critical task, the user is assured that the most critical questions are asked first. Therefore, if questioning must be prematurely terminated, the time has been used to maximum advantage.

The current Merit system, as implemented in BATTLE, is theoretically based on an infinitesimal change to the value of a top node caused by an infinitesimal change to the value of a node on a tree. Merit is efficient, so that it can handle a BATTLE knowledge base with about 700 nodes on the inference network and a data base with about 5,000 data items. Although Merit has worked well experimentally, it is desirable to extend it to handle six situations.

1. Merit should be defined and used relative to any set of nodes, not just one top node.
2. Merit should handle a potential finite change to the value.
3. Merit should be based on a directed acyclic graph, not just a tree.
4. Merit should handle the propagation of distributions, not just single values.
5. Merit should be made practical for a larger knowledge base (7,000 nodes on the inference network).
6. Merit should be made practical for a larger data base (50,000 data items).

Finally, we are looking at ways in which expert systems interact with their environments. Within the context of solving classification problems, we intend to investigate two related aspects of this interaction: *"query generation"* and *explanation* . Some simple expert systems proceed from a fixed set of input data to an evaluation of that data. For most problem domains, however, the space of possibly relevant information is large, and some or all of this information may have costs associated with its acquisition. Thus, computational and other costs can be reduced by some mechanism that intelligently selects appropriate queries designed to solicit information that is relevant and cost effective in terms of the problem being solved. Expert systems for complex problem domains must also be able to generate explanations for their actions. Unless the system operates entirely autonomously, users must be apprised of the rational for system actions. There is a particular need for explanations tailored for system users rather than system designers.

In its general form, the query generation problem arises any time an expert system requires outside information. This information typically can be provided my one of three sources: the *user, sensors,* or *"other systems"* . Most current systems get information about a particular problem by directly querying a user. Data entry is slow and user patience limited, so it is important that the number of questions asked be restricted and the particular questions asked seem reasonable to the user. Some expert systems have direct access to sensors. The selection of sensor data becomes important if the number of sensors is large, sensor data can conflict, and/or costs are associated with sensing. Query generation may require substantial domain knowledge, particularly when different queries can interact. (Destructive testing provides an extreme example) At times, the set of possible queries is large and poorly structured, requiring that complex queries be synthesized rather than selected from a preexisting collection. Finally, expert systems are now being integrated into larger software systems. Such systems may need to access information in other components such as a large data base.

These queries are likely to require substantial computational resources, and should be chosen in an appropriate manner. In typical systems today, queries are chosen in either a depth first or a breadth first manner (relative to the deduction tree). Intelligent query generation will become increasingly more important as expert systems techniques are applied to increasingly complex problems.

*Explanation facilities*

Explanations are required by both system users and system designers. Users require a rationale for system actions in order to evaluate the utility of those actions. System designers need to be able to monitor and modify the reasoning processes used by the system. The simplest explanation facilities provide an exact trace of the line of reasoning used to come to a decision. (For example, in rule-based systems, an explanation may consist of an ordered list of rule firings) More sophisticated explanations are possible if the line of reasoning is reformulated by mapping it into a level of abstraction more easily comprehended. Major work still needs to be done in accounting for the different audiences requiring explanation capabilities. The heuristics needed to solve complicated problems often make the workings of an expert system difficult to understand, particularly by users not familiar with the internal structure of the system. A tradeoff exists between efficient and comprehensible inference. Comprehensibility can be improved by including in expert systems knowledge structures and inference mechanisms specific to explanation generation.

*Joint Industry/University Projects*

The researchers are involved in a number of projects which combine the research interests of the SOLVER project and the application interests of companies. These projects include knowledge acquisition and design in the following areas:

Computer system and hardware diagnosis Chemical reactor design Optimal experimental design for new product development and off line quality control.

This kind of joint research is promoted by the Microelectronics and Information Sciences Center of the University of Minnesota (MEIS).

Long range --

Our long range objective is to improve the methodology of the "knowledge capturing" process that occurs in the early stages of the development of expert systems when problem decomposition and solution strategies are being specified. Several related questions of interest include:

What are the performance consequences of different approaches, how can these consequences be evaluated, and what tools can assist in making the best choice?

How can organizations be determined which not only perform well, but are structured so as to facilitate knowledge acquisition from human experts?

We will continue to explore these questions in areas of design and management as well as in law, management, and medicine.

*B. Justification and Requirements for Continued SUMEX Use*

Our current model development takes advantage of the sophisticated Lisp programming environments on SUMEX and local facilities. Although much current work with Galen is done using a version running on a local VAX 11/780, we continue to benefit from the interaction with other researchers facilitated by the SUMEX system. We expect to use SUMEX to allow other groups access to the Galen program. We also plan to continue use of the knowledge engineering tools available on SUMEX.

We are working toward a Commonlisp implementation of the Galen system and expect to rely heavily on Commonlisp for future projects.

One of our students implemented a demonstration legal expert system in EMYCIN using the SUMEX resource, and we still find that the resource is valuable for making available major systems which we do not have locally, such as EMYCIN.

*C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM*

Our current grant from MEIS has permitted us to purchase four Perq 2 AI workstations for our Artificial Intelligence laboratory. The availability of Commonlisp on these machines is one reason why we expect to make use of that language in the future.

SUMEX will continue to be used for collaborative activities and for program development requiring tools not available locally.

*D. Recommendations for Future Community and Resource Development*

As a remote site, we particularly appreciate the communications that the SUMEX facility provides our researchers with other members of the community. We, too, are moving toward a workstation based development environment, but we hope that SUMEX will continue to serve as a focal point for the medical AI community. In addition to communication and sharing of programs, we are interested in development of Commonlisp based knowledge engineering tools. The continued existence of the SUMEX resource is very important to us.

# IV.C. Pilot Stanford Projects

Following are descriptions of the informal pilot projects currently using the Stanford portion of the SUMEX-AIM resource, pending funding, full review, and authorization.

In addition to the progress reports presented here, abstracts for each project are submitted on a separate Scientific Subproject Form.

# IV.C.1. REFEREE Project

**REFEREE Project**

**Bruce G. Buchanan, Ph.D., Principal Investigator**
**Computer Science Department**
**Stanford University**

Byron W. Brown, Ph.D., Co-Principal Investigator
Department of Medicine
Stanford University

Daniel E. Feldman, Ph.D., M.D., Associate Investigator
Department of Medicine
Stanford University

## I. SUMMARY OF RESEARCH PROGRAM

### *A. Project Rationale*

The goals of this project are related both to medical science and artificial intelligence: (a) use AI methods to allow the informed but non-expert reader of the medical literature to evaluate a randomized clinical trial, and (b) use the interpretation of the medical literature as a test problem for studies of knowledge acquisition and fusion of information from disparate sources. REFEREE will be used to evaluate the medical literature of clinical trials to determine the quality of a clinical trial, make judgements on the efficacy of the treatment proposed, and synthesize rules of clinical practice. The research is an initial step toward a more general goal - building computer systems to help the clinician and medical scientist read the medical literature more critically and more rapidly.

### *B. Medical Relevance*

The explosive growth of the medical literature has created a severe information gap for the busy clinician. Most physicians can afford neither the time required to study all the pertinent journal articles in their field, nor the risk of ignoring potentially significant discoveries. The majority of clinicians, in fact, have little sophistication in epidemiology and statistics; they must nonetheless base their pragmatic decisions on a combination of clinical experience and published literature. The clinician's computerized assistant must ferret out useful maxims of clinical practice from the medical literature, pass judgment on the quality of medical reports, evaluate the efficacy of proposed treatments, and adjudicate the interpretation of conflicting and even contradictory studies.

### *C. Highlights of Progress*

REFEREE, a rule-based system built upon the EMYCIN framework, partially encodes the epidemiological knowledge of two highly regarded experts at Stanford, a biostatistician (Dr. Bill Brown) and a clinician (Dr. Dan Feldman). The REFEREE system, in particular, allows the informed but non-expert reader of the medical literature to study the quality, believability, statistical precision, and overall merit of a randomized clinical trial. REFEREE structures its analysis into five categories: the CREDENTIALS of the author and the institution behind the study; the PLANNING of the trial, including stopping rules, selection of subgroups, and standards for evaluating the endpoints; the EXECUTION of the study, including the degree of blinding and the adherence to planning standards; the REPORTING of results; and the statistical ANALYSIS applied to those results.

In the future, REFEREE will alleviate the knowledge-acquisition bottleneck for an automated medical decision-maker: the program will evaluate the quality of a clinical trial, judge the efficacy of the treatment proposed therein, and synthesize rules of clinical practice. For the present, however, the fusion of knowledge from disparate sources remains a problem in pure AI. The system's designers, Dr. Bruce Buchanan and R. Martin Chavez, have instead focused their efforts on the refinement and deepening of REFEREE's epidemiological knowledge.

In its present incarnation, REFEREE holds its knowledge of controlled clinical methodology in a database of 250 rules. The investigators working on the REFEREE project have, over the last year, produced a prototype consultant that evaluates the design and reporting of a randomized control trial. Clearly, REFEREE must possess a great deal of statistical sophistication before it can usefully serve the critical reader of medical literature.

In the spring of 1986, Chavez undertook the task of structuring the vast corpus of statistical techniques used in clinical trials. The current version of REFEREE contains, in preliminary form, Prof. Brown's expert knowledge of biostatistics. REFEREE evaluates each statistical procedure described by the authors of the paper. The automated consultant then determines the most appropriate method for the problem at hand, based on the design of the trial and the hypotheses to be tested. REFEREE checks critical assumptions, looks for possible statistical abuses, verifies adjustments, and re-computes the statistics. In a beta-blocker study that employs the Cox proportional-hazards model, for instance, REFEREE will analyze the Kaplan-Meier survival curve and verify or reject the presence of a significant treatment effect.

In order to evaluate the paper's presentation of a statistical test, REFEREE must apply three kinds of knowledge. First, REFEREE examines the problem at hand and exhaustively establishes a list of relevant techniques. REFEREE applies rules of the following form in order to build the list of appropriate tests:

```
IF: 1) the data consist of two independent random samples, and
    2) A: outlying data prevent the application of a test based
          on sample means, or
       B: it is not known whether the data are from normally-
          distributed populations,
       and
    3) no censored data are present,
THEN:  The nonparametric ranktest of Wilcoxon is a reasonable test.
```

REFEREE presently possesses about fifty rules that define the range of admissible tests.

Second, REFEREE must know how to perform each kind of statistical test in order to verify the paper's results. Algorithmic knowledge (expressed as Interlisp code), for instance, enables REFEREE to calculate chi-squared for JxK contingency tables, evaluate the binomial distribution, and compute the Gehan rank statistic for survival curves. REFEREE knows how to apply the majority of statistical methods in common usage.

Finally, REFEREE must know how to test hypotheses. In order to look for significance in a contingency table, for instance, REFEREE must know how to compute P values and critical values from the chi-squared distributions. REFEREE can understand and interpret all the major distributions used in clinical epidemiology.